

1 *Biometrika* (2008), **xx**, x, pp. 1–13
2 © 2007 Biometrika Trust
3 Printed in Great Britain

4 **Approximate Bayesian computation (ABC) gives exact results 5 under the assumption of model error**

6 **BY RICHARD D. WILKINSON**

7 *Department of Probability and Statistics, University of Sheffield, S3 7RH, United Kingdom*

8 r.d.wilkinson@sheffield.ac.uk

9 **SUMMARY**

10 Approximate Bayesian computation (ABC) or likelihood-free inference algorithms are used
11 to find approximations to posterior distributions without making explicit use of the likelihood
12 function, depending instead on simulation of sample data sets from the model. In this paper we
13 show that under the assumption of the existence of a uniform additive model error term, ABC
14 algorithms give exact results when sufficient summaries are used. This interpretation allows the
15 approximation made in many previous application papers to be understood, and should guide the
16 choice of metric and tolerance in future work. ABC algorithms can be generalized by replacing
17 the 0-1 cut-off with an acceptance probability that varies with the distance of the simulated data
18 from the observed data. The acceptance density gives the distribution of the error term, enabling
19 the uniform error usually used to be replaced by a general distribution. This generalization can
20 also be applied to approximate Markov chain Monte Carlo algorithms. In light of this work,
21 ABC algorithms can be seen as calibration techniques for implicit stochastic models, inferring
22 parameter values in light of the computer model, data, prior beliefs about the parameter values,
23 and any measurement or model errors.

24 *Some key words:* Approximate Bayesian computation; calibration; implicit inference; likelihood-free inference; Monte
25 Carlo.

26 **1. INTRODUCTION**

27 Approximate Bayesian computation (ABC) algorithms are a group of methods for performing
28 Bayesian inference without the need for explicit evaluation of the model likelihood function
29 (Beaumont et al. (2002); Marjoram et al. (2003); Sisson et al. (2007)). The algorithms can be
30 used with implicit computer models (Diggle & Gratton (1984)) that generate sample data sets
31 rather than return likelihoods. ABC methods have become popular in the biological sciences
32 with applications in genetics (see, for example, Siegmund et al. (2008); Foll et al. (2008)), epi-
33 demiology (Blum & Tran (2008); Tanaka et al. (2006)) and population biology (Ratmann et al.
34 (2007); Hamilton et al. (2005); Cornuet et al. (2008)) most common. This popularity is primarily
35 due to the fact that the likelihood function, which can be difficult or impossible to compute for
36 some models, is not needed in order to do inference. However, despite their popularity little is
37 known about the quality of the approximation they provide beyond results shown in simulation
38 studies.

39 In this paper we give a framework in which the accuracy of ABC methods can be understood.
40 The notation throughout this paper is as follows. Let θ denote the vector of unknown model
41 parameters we wish to infer, and let $\eta(\cdot)$ denote the computer model. We assume $\eta(\cdot)$ is stochastic,
42 so that the model repeatedly run at θ will give a range of possible model outputs, and write

49 $X \sim \eta(\theta)$ to denote that X has the same distribution as the model run at θ . To distinguish the
 50 model output from the observed data, let D denote the observations. The aim is to calibrate the
 51 model to the data, in order to learn about the true value of the parameter. The Bayesian approach
 52 is to find the posterior distribution of θ given D , given by

$$53 \quad 54 \quad 55 \quad \pi(\theta | D) = \frac{\pi(D | \theta)\pi(\theta)}{\pi(D)}.$$

56 Throughout this paper, $\pi(\cdot)$ is used to denote different probability densities, and $\pi(\cdot | \cdot)$ conditional
 57 densities, with the context clear from the arguments. Above, $\pi(\theta)$ is the prior distribution,
 58 $\pi(D | \theta)$ is the likelihood of the data under the model given parameter θ (the probability distribution
 59 of $\eta(\theta)$), $\pi(\theta | D)$ is the posterior distribution, and $\pi(D)$ is the evidence for the model.

60 It is usual in Bayesian inference to find that the normalizing constant $\pi(D)$ is intractable,
 61 and a wide range of Monte Carlo techniques have been developed to deal with this case (Liu
 62 (2001)). Doubly-intractable distributions are distributions which have a likelihood function
 63 $\pi(D | \theta) = q(D | \theta)/c(\theta)$ which is known only up to a normalizing constant, $c(\theta)$, which is in-
 64 tractable. Standard Monte Carlo techniques do not apply to these distributions, and Murray et al.
 65 (2006) have developed algorithms which can be used in this case. ABC methods are Monte
 66 Carlo techniques developed for use with completely-intractable distributions, where the likeli-
 67 hood function $\pi(D | \theta)$ is not even known up to a normalizing constant. ABC algorithms, some-
 68 times called likelihood-free algorithms, enable inference using only simulations generated from
 69 the model, and do not require any evaluation of the likelihood. The most basic form of the ABC
 70 algorithm is based on the rejection algorithm, and is as follows:

71 **Algorithm A: approximate rejection algorithm**

72 A1 Draw $\theta \sim \pi(\theta)$
 73 A2 Simulate X from the model $X \sim \eta(\theta)$
 74 A3 Accept θ if $\rho(X, D) \leq \delta$.

75 Here, $\rho(\cdot, \cdot)$ is a distance measure on the model output space, and δ is a tolerance determining
 76 the accuracy of the algorithm. Accepted values of θ are not from the true posterior distribution,
 77 but from an approximation to it, written $\pi(\theta | \rho(D, X) \leq \delta)$. When $\delta = 0$ this algorithm is exact
 78 and gives draws from the posterior distribution $\pi(\theta | D)$, whereas as $\delta \rightarrow \infty$ the algorithm gives
 79 draws from the prior. While smaller values of δ lead to samples which better approximate the
 80 true posterior, they also lead to lower acceptance rates in step A3 than larger values, and so
 81 more computation must be done to get a given sample size. Consequently, the tolerance δ can be
 82 considered as controlling a trade-off between computability and accuracy.

83 Several extensions have been made to the approximate rejection algorithm. If the data are high
 84 dimensional, then a standard change to the algorithm is to summarize the model output and data,
 85 using a summary statistic $S(\cdot)$ to project X and D onto a lower dimensional space. Algorithm A
 86 is then changed so that step A3 reads

87 A3' Accept θ if $\rho(S(X), S(D)) \leq \delta$.

88 Ideally, $S(\cdot)$ should be chosen to be a sufficient statistic for θ . However, if the likelihood is
 89 unknown, then sufficient statistics cannot be identified. Summarizing the data and model output
 90 using a non-sufficient summary adds another layer of approximation on top of that added by the
 91 use of the distance measure and tolerance, but again, it is not known what effect any given choice
 92 for $S(\cdot)$ has on the approximation.

97 Beaumont et al. (2002) contains two innovations; they replace the discrete 0-1 cut-off in step
 98 A3 with a weighting scheme, using an Epanechnikov kernel to weight each value of θ according
 99 to the value of the metric $\rho(D, X)$, so that large weights are assigned to values of θ which pro-
 100 duce model output close to the measurement and small weights to those values which produce
 101 output distant from D . They then use the weighted sample (θ_i, w_i) to train a local-linear regres-
 102 sion to model the posterior density. Blum & Fran ois (2008) have since extended this by using an
 103 adaptive heteroscedastic model to estimate the posterior from the weighted sample. While it has
 104 been shown that using a weighting scheme improves the accuracy of the approximate rejection
 105 algorithm in several situations, it is still unclear what the approximation represents or why any
 106 given weighting should be preferred over any other.

107 In this paper it is shown that the basic approximate rejection algorithm can be interpreted as
 108 performing exact inference in the presence of uniform model or measurement error. Similarly, the
 109 weighting scheme used in Beaumont et al. (2002) corresponds to an error with an Epanechnikov
 110 distribution. In other words, it is shown that ABC gives exact inference for the wrong model,
 111 and we give a distribution for the model error term for whatever choice of metric and tolerance
 112 are used. This interpretation allows us to show the effect a given choice of metric, tolerance
 113 and weighting have had in previous applications, and should provide guidance when choosing
 114 metrics and weightings in future work. It is also shown that Algorithm A can be generalized
 115 to give inference under the assumption of a completely flexible form for the model error. We
 116 discuss how to model the model error, and show how some models can be rewritten to give exact
 117 inference.

118 Finally, ABC has been extended by Marjoram et al. (2003) from the rejection algorithm
 119 to approximate Markov chain Monte Carlo algorithms, and by Sisson et al. (2007) and
 120 Beaumont et al. (2008) to approximate sequential Monte Carlo algorithms. We extend the ap-
 121 proximate Markov chain Monte Carlo algorithm to give inference for a general form of error,
 122 and suggest methods for calculating Bayes factors and integrals for completely-intractable dis-
 123 tributions.

2. INTERPRETING ABC

124 In this section a framework is described which enables the effect a given metric and weighting
 125 have in ABC algorithms to be understood. This will then allow us to improve the inference by
 126 carefully choosing a metric and weighting which more closely represents our true beliefs. The
 127 key idea is to assume that there is a discrepancy between the best possible model prediction
 128 and the data. This discrepancy represents either measurement error on the data, or model error
 129 describing our statistical beliefs about where the model is wrong. George Box famously wrote
 130 that ‘all models are wrong’, and in order to link models to reality it is necessary to account for
 131 this model error when performing inference. In the context of deterministic models, this practice
 132 is well established (Campbell (2006); Goldstein & Rougier (2008); Higdon et al. (2008)), and
 133 should also be undertaken when linking stochastic models to reality, despite the fact that the
 134 variability in the model can seemingly explain the data as they are.

135 The framework introduced here uses the best input approach, similar to that given in
 136 Kennedy & O’Hagan (2001). We assume that the measurement D can be considered as a re-
 137 alization of the model run at its best input value, $\hat{\theta}$, plus an independent error term ϵ

$$138 \quad D = \eta(\hat{\theta}) + \epsilon. \quad (1)$$

139 The error ϵ might represent measurement error on D , or model error in $\eta(\cdot)$, or both, in which
 140 case we write $\epsilon = \epsilon_1 + \epsilon_2$. Discussion about the validity of Equation (1), and what ϵ represents

and how to model it are delayed until Section 3, and for the time being we simply consider ϵ to have density $\pi_\epsilon(\cdot)$. The aim is to describe our posterior beliefs about the best input $\hat{\theta}$ in light of the error ϵ , the data D , and prior beliefs about $\hat{\theta}$. Consider the following algorithm:

Algorithm B: probabilistic approximate rejection algorithm

- B1 Draw $\theta \sim \pi(\theta)$
- B2 Simulate X from the model $X \sim \eta(\theta)$
- B3 Accept θ with probability $\frac{\pi_\epsilon(D-X)}{c}$.

Here, c is a constant chosen to guarantee that $\pi_\epsilon(D - X)/c$ defines a probability. For most cases we will expect ϵ to have a modal value of 0, and so taking $c = \pi_\epsilon(0)$ will make the algorithm valid and also ensure efficiency by maximizing the acceptance rate. If D and X are both real valued arrays of matching dimension, then $D - X$ is simply the arithmetic pairwise difference. However, if D and X are not real-valued, for example, D and X could both be gene sequences, then $D - X$ represents the difference between the two data sets (we could write $D - X = \rho(D, X)$) and $\pi_\epsilon(\cdot)$ is a distribution on this space of differences.

The main innovation in this paper is to show that Algorithm B gives exact inference for the statistical model described above by Equation (1). This is essentially saying that ABC gives exact inference, but for the wrong model.

THEOREM 1. *Algorithm B gives draws from the posterior distribution $\pi(\hat{\theta} | D)$ under the assumption that $D = \eta(\hat{\theta}) + \epsilon$ and $\epsilon \sim \pi_\epsilon(\cdot)$ independently of $\eta(\hat{\theta})$.*

Proof. Let

$$I = \begin{cases} 1 & \text{if } \theta \text{ is accepted} \\ 0 & \text{otherwise.} \end{cases}$$

We then find that

$$\begin{aligned} \text{pr}(I = 1 | \theta) &= \int \text{pr}(I = 1 | \eta(\theta) = x, \theta) \pi(x | \theta) dx \\ &= \int \frac{\pi_\epsilon(D - x)}{c} \pi(x | \theta) dx. \end{aligned}$$

This gives that the distribution of accepted values of θ is

$$\pi(\theta | I = 1) = \frac{\pi(\theta) \int \pi_\epsilon(D - x) \pi(x | \theta) dx}{\int \pi(\theta') \int \pi_\epsilon(D - x) \pi(x | \theta') dx d\theta'}.$$

To complete the proof we must find the posterior distribution of the best model input $\hat{\theta}$ given the data D under the assumption of model error. Note that $\pi(D | \eta(\hat{\theta}) = x) = \pi_\epsilon(D - x)$ which implies that the likelihood of θ is

$$\begin{aligned} \pi(D | \hat{\theta}) &= \int \pi(D | \eta(\hat{\theta}) = x, \hat{\theta}) \pi(x | \hat{\theta}) dx \\ &= \int \pi_\epsilon(D - x) \pi(x | \hat{\theta}) dx. \end{aligned}$$

193 Consequently, the posterior distribution of $\hat{\theta}$ is
 194

$$195 \pi(\hat{\theta} | D) = \frac{\pi(\hat{\theta}) \int \pi_\epsilon(D - x) \pi(x | \hat{\theta}) dx}{\int \pi(\theta) \int \pi_\epsilon(D - x) \pi(x | \theta) dx d\theta}$$

196 which matches the distribution of accepted values from Algorithm B. \square
 197

199 To illustrate the algorithm, we consider the toy example used in Sisson et al. (2007) and again
 200 in Beaumont et al. (2008) where analytic expressions can be calculated for the approximations.
 201

202 *Example 1.* Assume the model is a mixture of two normal distributions with a uniform prior
 203 for the mean:
 204

$$205 \eta(\theta) \sim \frac{1}{2} \mathcal{N}(\theta, 1) + \frac{1}{2} \mathcal{N}(\theta, \frac{1}{100}), \theta \sim \mathcal{U}[-10, 10].$$

207 Further assume that we observe $D = 0$, but that there is measurement error ϵ on this data. If
 208 $\epsilon \sim \mathcal{U}[-\delta, \delta]$, which is the assumption made when using Algorithm A with $\rho(x, 0) = |x|$, then it
 209 is possible to show that the approximation is

$$210 \pi(\theta | \epsilon \sim \mathcal{U}[-\delta, \delta], D = 0) \propto \Phi(\delta - \theta) - \Phi(-\delta - \theta) + \Phi(10(\epsilon - \theta)) - \Phi(-10(\epsilon + \theta))$$

211 for $\theta \in [-10, 10]$, where $\Phi(\cdot)$ is the standard Gaussian distribution function. Note that this is
 212 slightly different to the distribution given in Beaumont et al. (2008). An alternative to assuming
 213 uniform error, is to suppose that the error has a normal distribution $\epsilon \sim \mathcal{N}(0, \delta^2/3)$. It can then
 214 be shown that the posterior distribution of θ is
 215

$$216 \pi(\theta | \epsilon \sim \mathcal{N}(0, \frac{\delta^2}{3}), D = 0) \propto \frac{1}{2} \phi(\theta; 0, 1 + \frac{\delta^2}{3}) + \frac{1}{2} \phi(\theta; 0, \frac{1}{100} + \frac{\delta^2}{3})$$

217 truncated onto $[-10, 10]$. This is the approximation found when using Algorithm B with a Gaussian
 218 acceptance kernel, where $\phi(\cdot; \mu, \sigma^2)$ is the probability density function of a Gaussian distribution
 219 with mean μ and variance σ^2 . The value of the variance, $\delta^2/3$, is chosen to be equal
 220 to the variance of a $\mathcal{U}[-\delta, \delta]$ random variable. For large values of the tolerance δ , the difference
 221 between the two approximations can be significant (see Figure 1), but in the limit as δ tends to
 222 zero, the two approximations will be the same, corresponding to zero error.
 223

225 226 227 3. MODEL DISCREPANCY

228 The interpretation of ABC given by Theorem 1 allows us to revisit previous analyses done
 229 using the ABC algorithm, and to understand the approximation in the posterior in terms of the
 230 distribution implicitly assumed for the error term. If the approximate rejection algorithm (Algo-
 231 rithm A) was used to do the analysis, we can see that this is equivalent to using the acceptance
 232 probability
 233

$$234 \frac{\pi_\epsilon(r)}{c} = \begin{cases} 1 & \text{if } \rho(r) \leq \delta \\ 235 0 & \text{otherwise} \end{cases}$$

236 where r is the distance between the simulated and observed data. This says that Algorithm A
 237 gives exact inference for the model which assumes a uniform measurement error on the region
 238 defined by the 0-1 cut-off, i.e.,
 239

$$240 \epsilon \sim \mathcal{U}\{x : \rho(x, D) \leq \delta\}.$$

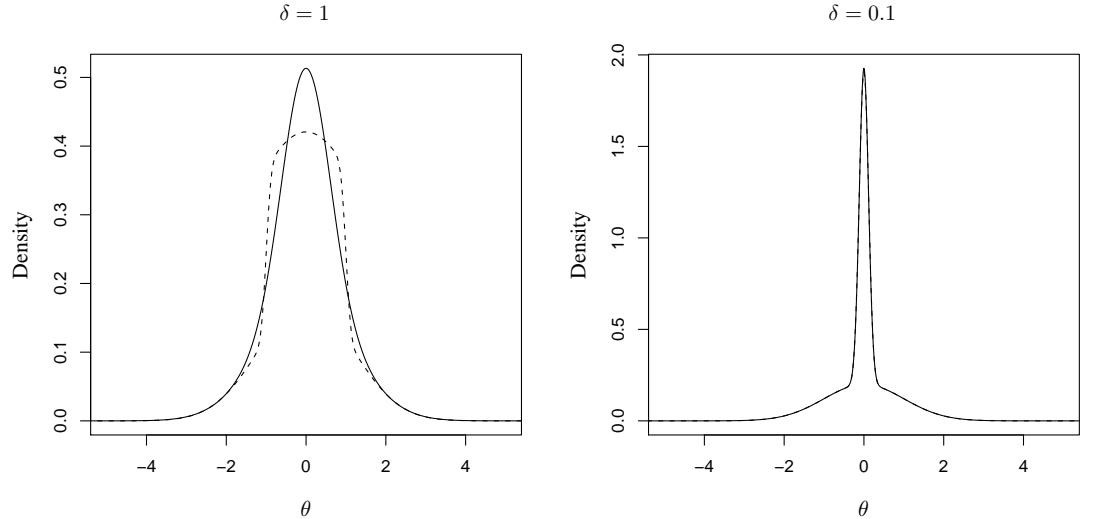


Fig. 1. The posterior distributions found when using Algorithm A (solid line) and Algorithm B (dashed line) with a Gaussian acceptance kernel. The left plot is for $\delta = 1$ and the right plot for $\delta = 0.1$.

If $\rho(\cdot, \cdot)$ is a Euclidean metric, $\rho(D, x) = (x - D)^T(x - D)$, this is equivalent to assuming uniform measurement error on a ball of radius δ about D .

The weighting scheme used in Beaumont et al. (2002), and since used in numerous application papers, used an Epanechnikov kernel $K_\delta(\|x - D\|)$ to weight each value of θ . The form of this density is

$$K_\delta(r) = \frac{3}{4\delta} \left(1 - \frac{r^2}{\delta^2}\right) \mathbb{I}_{r \leq \delta}. \quad (2)$$

The kernel gives the probability density function of the error term, and thus analysis done using an Epanechnikov kernel give results which assume error with this distribution. In most situations, it is likely to be a poor choice for a model of the measurement error, because the tails of the distribution are short, with zero mass outside of the interval $[-\delta, \delta]$.

There are two ways we can choose to view the error term; either as measurement error or model error. Interpreting ϵ to represent measurement error is relatively straight forward, as scientists usually hold beliefs about the distribution and magnitude of measurement error on their data. For most problems, assumptions of uniform measurement error will be inappropriate, and so using Algorithm A with a 0-1 cut-off will be inappropriate. But we have shown how to replace this uniform assumption with a distribution which more closely represents the beliefs of the scientist. Although the distribution of the measurement error will often be completely specified by the scientist, for example zero-mean Gaussian error with known variance, it is possible to include unknown parameters for the distribution of ϵ in θ and infer them along with the model parameters. Care needs to be taken to choose the constant c so that the acceptance rate in step B3 is less than one for all values of the parameter, but other than this it is in theory simple to infer error parameters along with the model parameters. So for example, if $\epsilon \sim \mathcal{N}(0, \sigma^2)$, where σ^2 is unknown, we could include σ^2 in θ .

289 Some models have sampling or measurement error built into the computer code so that the
 290 model output includes a realization of this noise. Rather than coding the noise process into the
 291 model, it will sometimes be possible to rewrite the model so that it outputs the latent underlying
 292 signal. If the likelihood of the data given the latent signal is computable (as it often is), then it
 293 may be possible to analytically account for the noise with the acceptance probability $\pi_\epsilon(\cdot)$. ABC
 294 methods have proven most popular in fields such as genetics, epidemiology, and population bi-
 295 ology, where a common occurrence is to have data generated by sampling a hidden underlying
 296 tree structure. In many cases, it is the partially observed tree structure which causes the likeli-
 297 hood to be intractable, and given the underlying tree the sampling process will have a known
 298 distribution. If this is the case (and if computational constraints allow), we can use the proba-
 299 bilistic ABC algorithm to do the sampling to give exact inference without any assumption of
 300 model error. Note that if the sampling process gives continuous data, then exact inference using
 301 the rejection algorithm would not be possible, and so this approach has the potential to give a
 302 significant improvement over current methods.

303 *Example 2.* To illustrate the idea of rewriting the model in order to do analytic sampling,
 304 we describe a version of the problem considered in Plagnol & Tavaré (2004). Their aim was to
 305 use the primate fossil record to date the divergence time of the primates. They used an inho-
 306 mogeneous branching process to model speciation, with trees rooted at time $t = \tau$, and simu-
 307 lated forwards in time to time $t = 0$, so that the depth of the tree, τ , represents the divergence
 308 time of interest. The branching process is parametrized by λ , which can either be estimated and
 309 fixed, or treated as unknown and given a prior distribution. Time is split into geologic epochs
 310 $\tau < t_k < \dots < t_1 < 0$, and the data consist of counts of the number of primate species that have
 311 been found in each epoch of the fossil record, $D = (D_1, \dots, D_k)$. Fossil finds are modelled by
 312 a discrete marking process on the tree, with each species having equal probability α of being
 313 preserved as a fossil in the record. If we let N_i be the cumulative number of branches that ex-
 314 ist during any point of epoch i , then the model used for the fossil finds process can be written
 315 as $D_i \sim \text{Binomial}(N_i, \alpha)$. The distribution of $N = (N_1, \dots, N_{14})$ cannot be calculated explic-
 316 itely and so we cannot use a likelihood based approach to find the posterior distribution of the
 317 unknown parameter $\theta = (\lambda, \tau, \alpha)$. The ABC approach used in Plagnol & Tavaré (2004) was to
 318 draw a value of θ from its prior, simulate a sample tree and fossil finds, and then count the num-
 319 ber of simulated fossils in each epoch to find a simulated value of the data X . They then accepted
 320 θ if $\rho(D, X) \leq \delta$ for some metric $\rho(\cdot, \cdot)$ and tolerance δ . This gives an approximation to the pos-
 321 terior distribution of the parameter given the data and the model, where the approximation can
 322 be viewed as model or measurement error.

323 However, instead of approximating the posterior, it is possible in theory to rewrite the model
 324 and perform the sampling analytically to find the exact posterior distribution:

- 326 1. Draw $\theta = (\lambda, p, \alpha) \sim \pi(\cdot)$
- 327 2. Simulate a tree T using parameter λ and count N
- 328 3. Accept θ with probability $\prod_{i=1}^k \frac{N_i}{D_i} \alpha_i^{D_i} (1 - \alpha_i)^{N_i - D_i}$.

329 This algorithm gives exact draws from the posterior distribution of θ given D , and in theory
 330 there is no need for any assumption of measurement error. Note that θ can include parameter α
 331 for the sampling rate, to be inferred along with the other model parameters. However, this makes
 332 finding a normalizing constant in step 3 difficult. Without a normalizing constant to increase the
 333 acceptance rate, applying this algorithm directly will be slow for many values of D and k (the
 334 choice of prior distribution and number of parameters we choose to include in θ can also have
 335 a significant effect on the efficiency). A practical solution would be to add an error term and
 336

assume the presence of measurement error on the data (which is likely to exist in this case), in order to increase the acceptance probability in step 3. Approaching the problem in this way, it is possible to carefully model the error on D and improve the estimate of the divergence time.

Using ϵ to represent measurement error is straight forward, whereas using ϵ to model the model discrepancy (to account for the fact the model is wrong) is harder to conceptualize and not as commonly used. For deterministic models, the idea of using a model error term when doing calibration or data assimilation is well established and described for a Bayesian framework in Kennedy & O'Hagan (2001). They assume that the model run at its 'best' input, $\eta(\hat{\theta})$, is sufficient for the model when estimating $\hat{\theta}$. In other words, knowledge of the model run at its best input provides all the available information about the system for the purpose of prediction. If this is the case, then we can define ϵ to be the difference between $\eta(\hat{\theta})$ and D , and assume ϵ is independent of $\eta(\hat{\theta})$. Note that the error is the difference between the data and the model run at its best input, and does not depend on θ . If we do not include an error term ϵ , then the best input is the value of θ that best explains the data, given the model. When we include an error term which is carefully modelled and represents our beliefs about the discrepancy between $\eta(\cdot)$ and reality, then it can be argued that $\hat{\theta}$ represents the 'true' value of θ , and that $\pi(\hat{\theta} | D, \epsilon \sim \pi_\epsilon(\cdot))$ should be our posterior distribution for $\hat{\theta}$ in light of the data and the model.

For deterministic models, Goldstein & Rougier (2008) provide a framework to help think about the model discrepancy. To specify the distribution of ϵ , it can help to break the discrepancy down into various parts: physical processes not modelled, errors in the specification of the model, imperfect implementation etc. So for example, if $\eta(\cdot)$ represents a global climate model predicting average temperatures, then common model errors could be not including processes such as clouds, CO₂ emissions from vegetation etc., error in the specification might be using an unduly simple model of economic activity, and imperfect implementation would include using grid cells too large to accurately solve the underlying differential equations. In some cases it may be necessary to consider model and measurement error, $\epsilon + e$ say, and model each process separately. For stochastic models, as far as we are aware, no guidance exists about how to model the error, and indeed it is not clear what ϵ should represent.

To complicate matters further, for many models the dimension of D and X will be large, making it likely that the acceptance rate $\pi_\epsilon(X - D)$ will be small. As noted above, in this case it is necessary to summarize the model output and the data using a multidimensional summary $S(\cdot)$. Using a summary means that rather than approximating $\pi(\theta | D)$, the algorithms approximate $\pi(\theta | S(D))$. The interpretation of ϵ as model or measurement error still holds, but now the error is on the measurement $S(D)$ or the model prediction $S(X)$. If each element of $S(\cdot)$ has an interpretation in terms of a physical process, this may make it easier to break the error down into independent components. For example, suppose that we use $S(x) = (\bar{x}, s_{xx})$, the sample mean and variance of X , and that we then use the following acceptance density

$$\pi_\epsilon(S(X) - S(D)) = \pi_1(\bar{X} - \bar{D})\pi_2(s_{XX} - s_{DD}).$$

This is equivalent to assuming that there are two sources of model error. Firstly, the mean prediction is assumed to be wrong, with the error distributed with density $\pi_1(\cdot)$. Secondly, it assumes that the model prediction of the variance is wrong, with the error having distribution $\pi_2(\cdot)$. It also assumes that the error in the mean prediction is independent of the error in the variance prediction. This independence is not necessary, but helps with visualization and elicitation. For this reason it can be helpful to choose the different components of $S(\cdot)$ so that they are close to independent (independence may also help increase the acceptance rate). Another possibility for choosing $S(\cdot)$ is to use principal component analysis (if $\dim(X)$ is large) to find a smaller

385 number of uncorrelated summaries of the data which may have meaningful interpretations. In
 386 general however, it is not known how to choose good summaries. Joyce & Marjoram (2008)
 387 have suggested a method for selecting between different summaries and for deciding how many
 388 summaries it is optimal to include. However, more work is required to find summaries which are
 389 informative, interpretable and for which we can describe the model error.

390 Finally, once we have specified a distribution for ϵ , we may find the acceptance rate is too small
 391 to be practicable and that it is necessary to compromise (as in Example 2 above). A pragmatic
 392 way to increase the acceptance rate is to use a more disperse distribution for ϵ . This moves us
 393 from the realm of using ϵ to model an error we believe exists, to using it to approximate the true
 394 posterior. This is currently how most ABC methods are used. However, even when making a
 395 pragmatic compromise, the interpretation of the approximation in terms of an error will allow us
 396 to think more carefully about how to choose between different compromise solutions.

397 *Example 3.* One of the first uses of an ABC algorithm was by Pritchard et al. (1999), who used
 398 a simple stochastic model to study the demographic history of the Y chromosome, and used an
 399 approximate rejection algorithm to infer mutation and demographic parameters for their model.
 400 Their data consisted of 445 Y chromosomes sampled at eight different loci from a mixture of
 401 populations from around the world, which they summarized by just three statistics: the mean
 402 (across loci) of the variance of repeat numbers V , the mean effective heterozygosity H , and
 403 the number of distinct haplotypes N . The observed value of the summaries for their sample
 404 was $D \equiv (V, H, N)^T = (1.149, 0.6358, 316)^T$. They elicited prior distributions for the mutation
 405 rates from the literature, and used diffuse priors for population parameters such as the growth rate
 406 and the effective number of ancestral Y chromosomes. Population growth was modelled using
 407 a standard coalescent model growing at an exponential rate from a constant ancestral level, and
 408 various different mutation models were used to simulate sample values for the three summaries
 409 measured in the data. They then applied Algorithm A using the metric
 410

$$411 \quad 412 \quad 413 \quad \rho(D, X) = \prod_{i=1}^3 \frac{D_i - X_i}{D_i} \quad (3)$$

414 where X is a triplet of simulated values for the three summaries statistics. They used a tolerance
 415 value of $\delta = 0.1$, which for their choice of metric corresponds to an error of 10% on each
 416 measurement. This gives results equivalent to assuming that there is independent uniform measurement
 417 error on the three data summaries, so that the true values of the three summaries have
 418 the following distributions

$$419 \quad 420 \quad V \sim \mathcal{U}[1.0341, 1.2624], \quad H \sim \mathcal{U}[0.58122, 0.71038], \quad N \sim \mathcal{U}[284, 348].$$

421 Beaumont et al. (2002) used the same model and data set to compare the relative performance of
 422 Algorithm A with an algorithm similar to Algorithm B, using an Epanechnikov density applied
 423 to the metric value (3) for the acceptance probability $\pi_\epsilon(\cdot)$. They set a value of δ (the cut-off in
 424 Algorithm A and the range of the support for ϵ in Algorithm B) by using a quantile P_δ of the
 425 empirical distribution function of simulated values of $\rho(D, X)$, i.e., $P_{0.01}$ means they accepted
 426 the 1% of model runs with values closest to D . They concluded that Algorithm B gives more
 427 accurate results than Algorithm A, meaning that the distribution found using Algorithm B is
 428 closer to the posterior found when assuming no measurement error ($\delta = 0$).

429 The conclusion that Algorithm B is preferable to Algorithm A for this model is perhaps not
 430 surprising in light of what we now know, as it was not taken into account that both algorithms
 431 used the same value of δ . For Algorithm A this corresponds to assuming a measurement error
 432 with variance $\delta^2/3$, whereas using acceptance probability (2) is equivalent to assuming a

433 measurement error with variance $\delta^2/5$. Therefore, using Algorithm B uses measurement error
 434 only 60% as variable as that assumed in Algorithm A, and so it is perhaps not surprising that
 435 Algorithm B gives more accurate results in this case.

438 4. APPROXIMATE MARKOV CHAIN MONTE CARLO

439 For problems which have a tightly constrained posterior distribution (relative to the prior),
 440 repeatedly drawing independent values of θ from its prior distribution in the rejection algorithm
 441 can be inefficient. For problems with a high dimensional θ this inefficiency is likely to make the
 442 application of a rejection type algorithm impracticable. The idea behind Markov chain Monte
 443 Carlo (MCMC) is to build a Markov chain on θ and correlate successive observations so that
 444 more time is spent in regions of high posterior probability. Most MCMC algorithms, such as the
 445 Metropolis-Hastings algorithm, depend on knowledge of the likelihood function which we have
 446 assumed is not known. Marjoram et al. (2003) give an approximate version of the Metropolis-
 447 Hastings algorithm, which approximates the acceptance probability by using simulated model
 448 output with a metric and a 0-1 cut-off to approximate the likelihood ratio. This, as before, is
 449 equivalent to assuming uniform error on a set defined by the metric and the tolerance. As above,
 450 this algorithm can be generalized from assuming uniform measurement error to an arbitrary
 451 error term. Below, are two different algorithms to perform MCMC for the model described by
 452 Equation (1). The difference between the two algorithms lies in the assumption about on which
 453 space we choose to construct the stationary distribution. In Algorithm C we consider the state
 454 variable to belong to the space of parameter values Θ , and construct a Markov chain $\{\theta_1, \theta_2, \dots\}$
 455 which obeys the following dynamics:

457 **Algorithm C: probabilistic approximate MCMC 1**

458 C1 At time t , propose a move from θ_t to θ' according to transition kernel $q(\theta_t, \theta')$.

459 C2 Simulate $X' \sim \eta(\theta')$.

460 C3 Set $\theta_{t+1} = \theta'$ with probability

$$462 \quad r(\theta_t, \theta' | X') = \frac{\pi_\epsilon(D - X')}{c} \min \left(1, \frac{q(\theta', \theta_t) \pi(\theta')}{q(\theta_t, \theta') \pi(\theta_t)} \right), \quad (4)$$

463 otherwise set $\theta_{t+1} = \theta_t$.

464 An alternative approach is to introduce the value of the simulated output as an auxiliary variable
 465 and construct the Markov chain on the space $\Theta \times \mathcal{X}$, where \mathcal{X} is the space of model outputs.

466 **Algorithm D: probabilistic approximate MCMC 2**

467 D1 At time t , propose a move from $\psi_t = (\theta_t, X_t)$ to $\psi' = (\theta', X')$ with θ' drawn from transition
 468 kernel $q(\theta_t, \theta')$, and X' simulated from the model using θ' :

$$469 \quad X' \sim \eta(\theta')$$

470 D2 Set $\psi_{t+1} = (\theta', X')$ with probability

$$471 \quad r((\theta_t, X_t), (\theta', X')) = \min \left(1, \frac{\pi_\epsilon(D - X') q(\theta', \theta_t) \pi(\theta')}{\pi_\epsilon(D - X_t) q(\theta_t, \theta') \pi(\theta_t)} \right), \quad (5)$$

472 otherwise set $\psi_{t+1} = \psi_t$.

481 *Proof of convergence.* To show that these Markov chains converge to the required posterior
 482 distribution, it is sufficient to show that the chains satisfy the detailed balance equations

483
$$p(s, t)\pi(t) = p(t, s)\pi(s) \quad \text{for all } s, t$$

485 where $p(\cdot, \cdot)$ is the transition kernel of the chain and $\pi(\cdot)$ the required stationary distribution.

486 For Algorithm C the transition kernel is the product of $q(\theta, \theta')$ and the acceptance rate. To
 487 calculate the acceptance rate, note that in Equation (4) the acceptance probability is conditioned
 488 upon knowledge of X' and so we must integrate out X' to find $r(\theta, \theta')$. This gives the transition
 489 kernel for the chain:

490
$$p(\theta, \theta') = q(\theta, \theta') \int_{\mathcal{X}} \frac{\pi_{\epsilon}(D - X')}{c} \min \left(1, \frac{q(\theta', \theta_t)\pi(\theta')}{q(\theta_t, \theta')\pi(\theta_t)} \right) \pi(X' | \theta) dX'.$$

493 The target stationary distribution is

494
$$\pi(D | \theta) = \int_{\mathcal{X}} \pi_{\epsilon}(D - X)\pi(X | \theta) dX.$$

495 It is then simple to show that the Markov chain described by Algorithm C satisfies the detailed
 496 balance equations.

497 For Algorithm D, the transition kernel is

500
$$p((\theta, X), (\theta, X')) = q(\theta, \theta')\pi(X' | \theta') \min \left(1, \frac{\pi_{\epsilon}(D - X')q(\theta', \theta_t)\pi(\theta')}{\pi_{\epsilon}(D - X_t)q(\theta_t, \theta')\pi(\theta_t)} \right). \quad (6)$$

501 The Markov chain in this case takes values on in $\Theta \times \mathcal{X}$ and the required stationary distribution
 502 is

503
$$\pi(\theta, X | D) = \frac{\pi_{\epsilon}(D - X)\pi(X | \theta)\pi(\theta)}{\pi(D)} \quad (7)$$

504 It can then be shown that Equations (6) and (7) satisfy the detailed balance equations. \square

505 While Algorithm C is more recognisable as a generalization of the approximate MCMC algo-
 506 rithm given in Marjoram et al. (2003), Algorithm D is likely to be more efficient in most cases.
 507 This is because the ratio of model error densities that occurs in acceptance rate (5) is likely to
 508 result in larger probabilities than those given by Equation (4) which simply has a $\pi_{\epsilon}(D - x)/c$
 509 term instead. Algorithm D also has the advantage of not requiring a normalizing constant.

5. EXTENSIONS

5.1. Monte Carlo integration

510 Suppose our aim is to calculate expectations of the form

511
$$E(f(\theta) | D) = \int f(\theta)\pi(\theta | D) d\theta$$

512 where the expectation is taken with respect to the posterior distribution of θ . The simplest way
 513 to approach this is to draw a sample of θ values $\{\theta_i\}_{i=1, \dots, n}$ from $\pi(\theta | D)$ using Algorithm B,
 514 C or D and then approximate using the sum $n^{-1} \sum f(\theta_i)$. However, a more stable estimate can
 515 be obtained by using draws from the prior weighted by $\pi_{\epsilon}(D - X_i)$ as in Algorithm B. For each
 516 θ drawn from the prior in step B1, assign it weight $w_i = \pi(D - X_i)$. Then an estimator of the

529 required expectation is

$$530 \quad \frac{\sum f(\theta_i)w_i}{\sum w_i}.$$

531

532 Note that all values of θ drawn from the prior are used in the sum; there is no rejection step. This
533 is a direct extension of the estimate given in Beaumont et al. (2002) which used Epanechnikov
534 kernels to weight each value of θ .

536

537

538 *5.2. Approximate model selection*

539 The theoretical acceptance rate from the rejection algorithm (Algorithm A with $\delta = 0$) is equal
540 to the model evidence $\pi(D)$. The evidence from different models can then be used to calculate
541 Bayes factors which can be used to perform model selection (Kass & Raftery (1995)). It is possi-
542 ble to approximate the value of $\pi(D)$ by using the acceptance rate from Algorithm A. By doing
543 this for two or more competing models, we can perform approximate model selection, although
544 in practice this approach can be unstable (Wilkinson (2007)). The estimate of $\pi(D)$ can be im-
545 proved and made interpretable by using the weighted estimate

$$546 \quad \frac{1}{n} \sum_{i=1}^n \frac{1}{m} \sum_{j=1}^m \pi_\epsilon(D - X_i^j)$$

547

548

549 where $X_i^1, \dots, X_i^m \sim \eta(\theta_i)$ and $\theta_1, \dots, \theta_n \sim \pi(\cdot)$. This gives a more stable estimate than sim-
550 plifying taking the acceptance rate, and also tends to the exact value (as $n, m \rightarrow \infty$) for the model
551 given by Equation (1).

553

554

555 **6. DISCUSSION**

556

557 It has been shown in this paper that approximate Bayesian computation algorithm can be
558 considered to give exact inference under the assumption of model error. However, this is only
559 part of the way towards a complete understanding of ABC algorithms. In the majority of the
560 application papers using ABC methods, summaries of the data and model output have been used
561 to reduce the dimension of the output. It cannot be known whether these summaries are sufficient
562 for the data, and so in most cases the use of summaries means that there is another layer of
563 approximation. While this work allows us to understand the error assumed on the measurement
564 of the summary, it says nothing about what effect using the summary rather than the complete
565 data has on the inference. More work is required to discover this, and to help guide the choice of
566 which summaries to use.

567 The use of a model error term when making inferences is important if one wants to move
568 from making statements about the model to statements about reality. There has currently been
569 only minimal work done on modelling the discrepancy term for stochastic models. One way to
570 approach this is to view the model as deterministic, outputting a density $\pi_\theta(x)$ for each value of
571 the input θ (many realizations of $\eta(\theta)$ would be needed to learn $\pi_\theta(x)$). The discrepancy term
572 ϵ can then be considered as representing the difference between $\pi_\theta(x)$ and the true variability
573 inherent in the physical system (at least the variability given the level we choose to model it
574 at). Finally, it should be possible to generalize approximate sequential Monte Carlo methods in
575 a similar way to that done for the approximate rejection and approximate Markov chain Monte
576 Carlo algorithms.

577
578
579
580
581
582
583
584
ACKNOWLEDGEMENTS

I would like to thank Professor Paul Blackwell for his suggestion that the metric in the ABC algorithm might have a probabilistic interpretation. I would also like to thank Professor Simon Tavaré, Professor Tony O'Hagan, and Dr Oakley for useful discussions on an early draft of the manuscript.

584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
REFERENCES

BEAUMONT, M. A., CORNUET, J. M., MARIN, J. M. & ROBERT, C. P. (2008). Adaptivity for ABC algorithms: the ABC-PMC scheme. *In submission*.

BEAUMONT, M. A., ZHANG, W. & BALDING, D. J. (2002). Approximate Bayesian computation in population genetics. *Genetics* **162**, 2025–2035.

BLUM, M. G. B. & FRANÇOIS, O. (2008). Highly tolerant likelihood-free Bayesian inference: an adaptive non-linear heteroscedastic model. *In submission, available as arXiv:0809.4178v1*.

BLUM, M. G. B. & TRAN, V. C. (2008). Approximate Bayesian computation for epidemiological models: application to the Cuban HIV-AIDS epidemic with contact-tracing and unobserved infectious population. *In submission, available as arXiv:0810.0896v1*.

CAMPBELL, K. (2006). Statistical calibration of computer simulations. *Reliab. Eng. Syst. Safe.* **91**, 1358–1363.

CORNUET, J. M., SANTOS, F., BEAUMONT, M. A., ROBERT, C. P., MARIN, J. M., BALDING, D. J., GUILLEMARD, T. & ESTOUP, A. (2008). Inferring population history with DIY ABC: a user-friendly approach to approximate Bayesian computation. *Bioinformatics*, to appear.

DIGGLE, P. J. & GRATTON, R. J. (1984). Monte Carlo methods of inference for implicit statistical models. *J. R. Statist. Soc. B* **46**, 193–227.

FOLL, M., BEAUMONT, M. A. & GAGGIOTTI, O. (2008). An approximate Bayesian computation approach to overcome biases that arise when using amplified fragment length polymorphism markers to study population structure. *Genetics* **179**, 927–939.

GOLDSTEIN, M. & ROUGIER, J. C. (2008). Reified bayesian modelling and inference for physical systems (with discussion). *J. Stat. Plan. Infer.*, to appear.

HAMILTON, G., Currat, M., RAY, N., HECKEL, G., BEAUMONT, M. A. & EXCOFFIER, L. (2005). Bayesian estimation of recent migration rates after a spatial expansion. *Genetics* **170**, 409–417.

HIGDON, D., GATTIKER, J., WILLIAMS, B. & RIGHTLEY, M. (2008). Computer model calibration using high-dimensional output. *J. Am. Statis. Assoc.* **103**, 570–583.

JOYCE, P. & MARJORAM, P. (2008). Approximately sufficient statistics and Bayesian computation. *Stat. Appl. Genet. Mo. B*, 7, article 26.

KASS, R. E. & RAFTERY, A. E. (1995). Bayes factors. *J. Am. Statis. Assoc.* **90**, 773–795.

KENNEDY, M. & O'HAGAN, A. (2001). Bayesian calibration of computer models (with discussion). *J. R. Statist. Soc. B* **63**, 425–464.

LIU, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer Series in Statistics. Springer.

MARJORAM, P., MOLITOR, J., PLAGNOL, V. & TAVARÉ, S. (2003). Markov chain Monte Carlo without likelihoods. *Proc. Natl. Acad. Sci. USA* **100**, 15324–15328.

MURRAY, I., GHAHRAMANI, Z. & MACKAY, D. J. C. (2006). MCMC for doubly-intractable distributions. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence*. UAI.

PLAGNOL, V. & TAVARÉ, S. (2004). Approximate Bayesian computation and MCMC. In *Proceedings of Monte Carlo and Quasi-Monte Carlo Methods 2002*, H. Niederreiter, ed. Springer-Verlag.

PRITCHARD, J. K., SEIELSTAD, M. T., PEREZ-LEZAUN, A. & FELDMAN, M. W. (1999). Population growth of human Y chromosomes: a study of Y chromosome microsatellites. *Mol. Biol. Evol.* **16**, 1791–1798.

RATMANN, O., JORGENSEN, O., HINKLEY, T., STUMPF, M., RICHARDSON, S. & WIUF, C. (2007). Using likelihood-free inference to compare evolutionary dynamics of the protein networks of *H. pylori* and *P. falciparum*. *PLoS Comput. Biol.* **3**, 2266–2276.

SIEGMUND, K. D., MARJORAM, P. & SHIBATA, D. (2008). Modeling DNA methylation in a population of cancer cells. *Stat. Appl. Genet. Mo. B*, 7, article 18.

SISSON, S. A., FAN, Y. & TANAKA, M. M. (2007). Sequential Monte Carlo without likelihoods. *Proc. Natl. Acad. Sci. USA* **104**, 1760–1765.

TANAKA, M. M., FRANCIS, A. R., LUCIANI, F. & SISSON, S. A. (2006). Using approximate Bayesian computation to estimate tuberculosis transmission parameters from genotype data. *Genetics* **173**, 1511–1520.

WILKINSON, R. D. (2007). *Bayesian inference of primate divergence times*. Ph.D. thesis, University of Cambridge.